

1. Theoretical and Practical Rationality

It is traditional to distinguish theoretical reasoning, which most directly affects beliefs, from practical reasoning, which most affects plans and intentions.

Theoretical reasoning is reasoned change in belief, or reasoned no-change in belief. Practical reasoning is reasoned change in plans and intentions, or reasoned no-change.

There are many similarities between theoretical and practical reasoning, but there are also important differences. One difference is that a certain sort of arbitrary choice is permitted in practical but not in theoretical reasoning. Given a choice among several equally acceptable things to do, it can be rational arbitrarily to choose one and irrational not to make this arbitrary choice. But, given a choice among several equally acceptable beliefs, it is never rational to arbitrarily select one to believe. Rationality requires a suspension of judgment in that case.

Someone who is unable to make arbitrary choices of things to do suffers from a serious defect in practical rationality, like Buridan's ass, stuck halfway between equally attractive equally delicious piles of hay. With belief it is just the opposite.

Consider also wishful thinking. It is theoretically unreasonable, but practically reasonable. One's goals, wishes, and desires are relevant to practical reasoning in a way in which they are not relevant to theoretical reasoning. A desire for more money can rationally influence one's decision to take a better paying job. Such a desire should not rationally influence one's conclusion, reached after having been interviewed for the job, that the interview went well and one will be offered the job. It is irrational to let one's desire to have done well at the interview influence one's conclusion about whether or not one did well. To believe that something is so merely because one wants it to be so is theoretically unreasonable, whereas to decide to try to make something so because one wants it to be so is reasonable practical thinking. Desires can rationally influence the conclusions of

practical reasoning in a way that they cannot rationally influence the conclusions of theoretical thinking.

On the other hand, one's goals are in other ways relevant to one's theoretical reasoning. In particular, one's goals may give provide a very strong reason to think about one thing rather than another. One's desires help to set the problems that one's reasoning is aimed at resolving. It is overly simple to say that one's desires cannot rationally affect what conclusions are legitimately reached in theoretical reasoning. One's desires can rationally affect one's theoretical conclusions by affecting what questions one uses theoretical reasoning to answer.

The point about wishful thinking then is this: *given what question one is using theoretical reasoning to answer*, one's desires cannot rationally affect what answer one reaches to that question, whereas in practical reasoning one's desires can rationally influence not just the questions one considers but also what practical answers one gives to those questions.

1.1 Legitimate practical reasons for beliefs

There are complications. Sometimes one has good practical reasons to answer certain questions in a certain way. There might be evidence that people who believe they will recover quickly from a certain illness are more likely to recover quickly than are people who are otherwise the same but do not believe they will recover quickly. A person with that illness could then have a practical reason to believe that he or she will recover quickly. Similarly, a contestant in a tennis tournament may have a practical reason to believe that he or she will win the tournament, if there is evidence that possession of such a belief improves one's playing. Or, to take a slightly different example, a salesperson might make more sales if he or she believed in the value of the item to be sold; so such a salesperson might have a practical reason to have such a belief.

There are other cases as well. One may have a practical reason of loyalty to believe that one's friend is not guilty of a crime with which he or she has been charged. One may want to have certain beliefs in order to fit in with the "in

crowd” or to obtain a job that will go only to someone with the relevant beliefs. In these and various other cases there may be good practical reasons to believe something.

This indicates that the difference between practical reasons and theoretical reasons is not just a matter of what they are reasons for—intentions versus beliefs. The difference has to do with the way in which reasons are reasons.

Some of the examples just discussed mention a reason to believe something that does not make it more likely that the belief is true. Such reasons are sometimes called (e.g., by Foley, 1987) nonepistemic reasons for belief, in contrast with the more usual epistemic reasons for belief that do make a belief more likely to be true. We will see below that not all practical reasons for belief are nonepistemic. Some practical reasons do make a belief more likely to be true.

1.2 Inference versus implication (again)

I now want to return to the point with which this chapter began, namely that issues about inference and reasoning need to be distinguished from issues about implication and consistency. Inference and reasoning are psychological processes leading to possible changes in belief (theoretical reasoning) or possible changes in plans and intentions (practical reasoning). Implication is more directly a relation among propositions. Certain propositions imply another proposition when and only when, if the former propositions are true, so is the latter proposition.

It is one thing to say “A, B, and C imply D.” It is quite another thing to say, “If you believe A, B, and C, you should or may infer D.” The first of these remarks is a remark about implication. The second is a remark about inference. The first says nothing special about belief or any other psychological state, unless one of A, B, or C has psychological content, nor does the first remark say anything normative about what anyone “should” or “may” do (Goldman, 1986).

The first remark can be true without the second being true. It may require considerable logical talent, even genius, to see that the implication holds. A

person without such genius or talent may believe A, B, and C without having any reason at all to believe D.

Furthermore, a person who believes A, B, and C and realizes that A, B, and C imply D may also believe for very good reason that D is false. Such a person may now have a reason to stop believing one of A, B, or C rather than a reason to believe D.

Moreover, even someone who believes A, B, and C, who realizes that A, B, and C imply D, and who has no reason to think that D is false, may have no reason to infer D. Such a person may be completely uninterested in whether D is true or false and no reason to be interested.

Many trivial things follow from one's beliefs without one having any reason to infer them. One has no reasons to clutter one's mind with trivialities just because they follow from other things one believes.

These and related examples indicate that the connection between inference and implication is fairly complex. We will say more about it below.

Similar remarks hold for consistency. Just as issues about implication have to be distinguished from issues about reasonable inference, issues about consistency have to be distinguished from issues about rationality and irrationality.

Consistency and inconsistency are in the first instance relations among propositions and only indirectly relations among propositional attitudes.

Propositions are consistent when and only when it is possible for them all to be true together. Propositions are inconsistent when and only when it is not possible for them all to be true together.

So, it is one thing to say that certain propositions are inconsistent with each other and quite another to say that it is irrational for someone to believe those propositions. The first remark, unlike the second, says nothing special about belief or other psychological states, nor does it say anything normative. So, the first remark can be true without the second being true. Suppose one believes each of the propositions. The inconsistency may have gone unnoticed and may

be very difficult to discover. One's failure to discover the contradiction need not indicate any irrationality.

And, even if one notices that the beliefs in question are inconsistent, one may still have reasons to continue to accept each and it may be quite unclear which should be given up. One may not have the time or the ability to work out which should be given up. One may have more urgent matters to attend to before resolving this inconsistency in one's beliefs. One may be very hungry and want to have lunch before solving this problem. In the meantime, it may very well be rational for one to continue to believe them all.

1.3 Ideal reasoners?

Reasoning is subject to resource limits of attention, memory, and time. So, it is not rational to fill your time inferring trivial consequences of your beliefs when you have more important things to attend to. And you cannot be charged with irrationality for having inconsistent beliefs where it would be costly to avoid the inconsistency.

Some theories of rationality (Stalnaker, 1984) abstract away from resource limits. Such theories of ideal rationality are concerned with an "ideally rational agent" whose beliefs are always consistent and closed under logical implication. Other theorists argue that such an idealization appears to confuse rationality, ideal or otherwise, with logical genius and even divinity! Furthermore, as we shall see, it is unclear how to relate such an "ideal" to actual finite human beings with their resource-limited rationality.

We have already seen that ordinary rationality requires neither deductive closure nor consistency. Ordinary rationality does not require deductive closure, because one is not always rational to believe something simply because it is implied by one's other beliefs. Rationality does not require consistency, because one can be rational even though there are undetected inconsistencies in one's beliefs, and because it is not always rational to respond to the discovery of inconsistency by stopping whatever else one might be doing in order to eliminate the inconsistency.

Now consider an ideal agent with no limitations on memory, attention span, or time, with instantaneous and cost-free computational abilities. It is not obvious whether such an agent would have a reason to infer all the trivial consequences of his or her beliefs. Although it would not cost anything for the agent to draw all those consequences, even all infinitely many of them, there would also be no need to draw any of those consequences in the absence of a reason to be interested in them, since the agent can effortlessly compute any consequence whenever it might be needed.

Could an ideal agent's beliefs be inconsistent? Suppose ordinary classical logic. Then, if the agent's beliefs were also deductively closed, the agent would then believe everything, because everything follows from inconsistency in classical logic.

This raises a question as to how the ideal agent could recover from inconsistency. Ordinary rational agents deal with momentary inconsistency all the time. One believes P but discovers Q, realizing that P and Q cannot both be true. One believes that something will happen and is surprised when it does not happen. For a moment, one has inconsistent beliefs. Normally one quickly recovers from the inconsistency by abandoning one of the initial beliefs.

But consider the implications of this sort of surprise for an ideal deductively closed agent. If the beliefs of such an agent were even momentarily inconsistent, the agent could never rationally recover, since there would be no trace in the agent's beliefs of how the agent had acquired the inconsistent beliefs. Since rational recovery from inconsistency can appeal only to present beliefs, and, since the deductively closed agent has exactly the same beliefs no matter how he or she got into inconsistency, there is no way in which the deductively closed agent could use temporal criteria in retreating from inconsistency—the agent would have to recover in exactly the same way, no matter where he or she had started.

It is therefore quite unclear how ideal rational agents might deal with ordinary surprise. Various possibilities suggest themselves, but we need not consider

them, because in what follows we will be directly concerned with real rather than ideal rational agents.

2. General Conservatism versus Special Foundations

Ordinary reasoning is conservative in the sense that one starts where one is, with one's current beliefs and intentions. Rational and reasonable change in view consists in trying to make improvements in one's initial position. One's initial beliefs and intentions have a privileged position in the sense that one begins with them rather than with nothing at all with some special privileged part of those beliefs and intentions. So, for example, one ordinarily continues to believe something that one starts out believing in the absence of a special reason to doubt it.

An alternative and radical conception of rationality going back to Descartes (1637) requires beliefs to be associated with reasons or justifications. Such justifications appeal to other beliefs, themselves to be associated with justifications, and so forth, until certain special foundational beliefs are reached that are self-justifying and need no further justification. Special foundational beliefs include beliefs about immediate experiences, such as headaches and perceptions, obvious logical and mathematical axioms, and other intuitively obvious truths. Rational change in view and rational no-change in view are to start always from evidence—those propositions that are evident. Then one is to accept only what can be justified from one's evidence basis, in this view.

Recent versions of special foundationalism (Foley, 1987; Alston, 1989, Chisholm, 1982) do not require foundational beliefs to be guaranteed to be true. In the absence of special challenges to them, they are justified, but their initial justified status might be overridden by special reasons to doubt them.

Supposing that there are two conflicting theories of reasoning, special foundationalism and general conservatism, each can be described using the terminology of the other theory as follows: the special foundations theory is

conservative about all foundational beliefs, but only foundational beliefs, whereas general conservatism treats all beliefs as foundational.

One problem for special foundationalism is to explain why the special foundational beliefs should have the sort of special status assigned to them in special foundationalism. What distinguishes foundational beliefs from other beliefs in such a way that justifies conservatism with respect to the foundational beliefs but not the other beliefs?

A second and potentially more serious problem is that people simply do not keep track of their reasons for their nonfoundational beliefs. This is a problem because, according to special foundationalism, if one does not associate a complete enough justification with a given nonfoundational belief, then it is not rational or reasonable for one to continue to believe it. This may undermine a great many of one's beliefs. Few people can remember their reasons for various of their beliefs about geography or history; does that make it unreasonable for them to continue to believe, for example, that Rome is in Italy or that the Battle of Hastings was in 1066? Furthermore, when beliefs are acquired on the basis of perception, one rarely continues to remember the perceptual evidence on which the beliefs were based. If that makes it irrational to believe, for example, that one saw John yesterday, the reasonable thing to do is to abandon almost everything one believes!

The issue between the two approaches to reasoning amounts to a question about the burden of proof or justification. According to special foundationalism, the burden of justification falls on continuing to believe something, at least for nonfoundational beliefs. Any nonfoundational belief requires special justification.

Foundational beliefs do not require special justification. For them, what requires justification is failing to continue to believe them. Sometimes there is a reason to abandon a foundational belief, but such abandonment requires such a special reason.

According to general conservatism, the burden of justification is always on changing beliefs or intentions. One starts with certain beliefs and intentions and any change in them requires some special reason. any sort of change in belief or intention requires special justification. Merely continuing to believe what one believes or intends requires no special justification in the absence of a special challenge to that belief or intention.

Clearly general conservatism fits better with ordinary thinking. Special foundationalism would imply that it is irrational or unreasonable for a typical person to continue to believe most of what he or she believes.

3. Induction and Deduction

Some authors draw a mistaken contrast between deductive and inductive reasoning. This is a mistake. Deduction and induction are not two kinds of reasoning; in fact, they are not two kinds of anything.

Deduction is concerned with certain relations among propositions, especially relations of implication and consistency. Induction is not concerned with those or any similar sort of relation. Induction is a kind of reasoning; but deduction is not a kind of reasoning.

Deductive logic is sometimes presented via a certain notion of “proof” or “argument.” A proof or argument in this sense has premises, intermediate steps, and a final conclusion. Each step must follow logically from prior steps in accordance with one or another specific rule, sometimes (misleadingly) called a “rule of inference.” Sometimes a proof or argument is taken to be an instance of “deductive reasoning.” Deductive reasoning in this sense is sometimes contrasted with “inductive reasoning,” which allegedly takes a similar form, with premises, intermediate steps, and final conclusions, but with the following difference: deductive steps are always truth preserving, whereas inductive steps are not.

This way of looking at deduction and induction is very misleading. For one thing, consider the reasoning that goes into the construction of a deductive proof or argument. Except in the simplest cases, the best strategy is not to expect to

start with the premises, figure out the first intermediate step of the proof, then the second, and so on until the conclusion is reached. Often it is useful to start from the proposition to be proved and work backward. It is also useful to consider what intermediate results might be useful.

In other words, the so-called deductive rules of inference are not rules that you follow in constructing the proof. They are rules that the proof must satisfy in order to be a proof!

In other words, there is a difference between reasoning about a proof, involving the construction of a proof that must satisfy certain rules, and reasoning that proceeds temporally in the same pattern as the proof in accordance with those rules. One does not reason deductively in the sense that one reasons in the pattern of a proof. One can reason about a deductive proof, just as one can reason about anything else. But one's reasoning is not well represented by anything like a proof or argument in the above sense.

1.4 Deduction

Deduction is not a kind of inference or reasoning, although one can reason about deductions. Deduction is implication. A deduction or proof or argument exhibits an implication by showing intermediate steps.

Logic, conceived as the theory of deduction, is not by itself a theory of reasoning. In other words, it is not by itself a theory about what to believe or intend. It is not a theory concerning reasoned change in view or reasoned no change in view.

It is true that deductions, proofs, and arguments do seem relevant to reasoning. It is not just that you sometimes reason about deductions in the way you reason about your finances or where to go on your summer trip. It is an interesting and nontrivial problem to say just how deductions are relevant to reasoning, a problem that is hidden from view by talk of deductive and inductive reasoning, as if it were obvious that some reasoning follows deductive principles.

It must be useful to construct deductions in some reasoning about ordinary matters, and not just when one is explicitly reasoning about deductions or proofs. But why should it be useful to construct deductions? What role do they play in reasoning?

Sometimes one accepts a conclusion because one has constructed a proof of that conclusion from other things one accepts. But there are other cases in which one constructs a proof of something one already accepts in order to see what assumptions might account for it. In such a case, the conclusion that one accepts might be a premise of the proof. The connection between proofs and reasoning is therefore complex.

1.5 *Induction*

The term “induction” is sometimes is sometimes restricted to “enumerative induction” in which a generalization is inferred from its instances. But the term “induction” is often used more widely so as to include inference to the best explanation of one’s evidence. What makes one hypothesis “better” than another for the purpose of inference to the best explanation is an issue that we must discuss below.

Philosophers sometimes discuss a “problem of induction” (Bonjour, 1992): How can one be justified in drawing a conclusion that is not guaranteed to be true by one’s premises? But it is unclear what the problem of induction is supposed to be. Premises of an argument are to be distinguished from the starting points in reasoning, as I have already observed. The conclusion of an argument is not to be identified with the conclusion of reasoning, in the sense of what you end up with or “conclude” as the result of your reasoning. Even when reasoning culminates in the construction of an argument, the conclusion of the *argument* may be something one started off believing, and the conclusion of one’s *reasoning* may be to accept something that is a premise of an explanatory argument constructed as the result of inference to the best explanation.

Clearly, it would be stupid, indeed highly irrational, not to engage in inductive reasoning. One would no longer be able to learn from experience. One would

have no basis for any expectations at all about the future, since one's evidence entirely concerns the past.

The "problem of induction" is a creation of confusion about induction and deduction, arising from the deductive model of inference. Again, it is important to see that there are not two kinds of reasoning, deductive and inductive.

Deduction has to do with implication and consistency and is only indirectly relevant to what one should believe.

4. Coherence

Everything one believes is at least potentially relevant to the conclusions one can reasonably draw. Rationality is a matter of one's overall view, including one's beliefs and one's intentions.

If it is reasonable to change one's view in a certain way, we might say that one's view would be more rationally "coherent" if changed in that way. We can describe principles of rationality as principles of rational coherence. Adopting this terminology, we can (following Pollock, 1974) distinguish two sorts of coherence, positive and negative.

Negative coherence is merely the absence of incoherence. Beliefs and intentions are incoherent to the extent that they are inconsistent with each other or clash in other ways. Incoherence is something to be avoided, if possible, although I have observed that it is not always possible to avoid incoherence. One's beliefs might be inconsistent without one's realizing that they are. And, even if one is aware of inconsistency, one may not know of a sufficiently easy way to get rid of it. Still, to the extent that one is aware of incoherence in one's view, one has a reason to modify one's view in order to get rid of the incoherence, if one can do so without too much expense.

Here, then, is one way that logic, in the sense of the theory of deductive implication, might be relevant to the theory of rationality, through providing an account of (one kind of) incoherence or inconsistency.

Positive coherence among one's beliefs and intentions exists to the extent that they are connected in ways that allow them to support each other. Relevant connections may involve explanations, generalizations, and implications.

In many cases new conclusions are accepted because they are part of explanations that serve to connect and integrate prior beliefs. A detective tries to find the best explanation of the various clues and other evidence. Scientific theories are accepted because of the way they allow us to explain the phenomena. Acceptance of someone's testimony may involve acceptance of certain explanations, namely that the speaker's testimony is the result of a desire to say what's true plus certain beliefs about what is true, where these beliefs are the result of the speaker's being in a position to know what is true. When a car fails to start, one may infer that the battery is low: that best explains the disappointing sounds that occur when the key is turned.

In some cases generalizations are accepted on the basis of a prior acceptance of certain instances. We might treat this as a special case of inference to the best explanation, supposing that the accepted generalizations explain their instances. But then we must recognize that this is a different sort of explanation from the causal or quasi-causal explanation involved in accounting for testimony or the failure of a car to start. A general correlation does not cause its instances.

Implication is also an important connector. Sometimes one does accept a new conclusion because it is implied by things one already accepts. This is a second way in which deductive logic, as a theory of deductive relations, can be relevant to reasoning. Deductive logic, so construed, is a theory of deductive implication and implication can be a coherence giving connection.

In trying to develop a theory of rational coherence (something that does not yet exist), we might try to reduce some of these coherence giving factors to others. For example, we might try to reduce all cases to explanatory coherence. But that is implausible for many cases in which a conclusion is accepted because it is implied by other beliefs. What is the relevant explanation? We might say that

the argument leading to the conclusion explains “why the conclusion is true.” But that seems to stretch the notion of explanation.

Another idea would be to try to reduce all coherence to that involved in implication. That has some plausibility for certain explanations. And strict generalizations are related to their instances by implication. Explanations in physics often seem to work via implication, leading to the “deductive nomological model” of explanation (Hempel, 1965). But not all explanations take this form. Many explanations appeal to “default” principles that hold only “other things being equal” or “normally.” Such explanations are not easily treated as deductively valid arguments.

5. Simplicity

In trying to explain some data, it is reasonable to consider a very limited range among the infinitely many logically possible explanations. The reasonable inquirer restricts attention to the set of relatively simple hypotheses that might account for most of the data.

This is not to say very much, since it amounts to using the term “simple” for whatever the relevant factors are that restrict rational attention to a certain few hypotheses. Furthermore, we are concerned with *relative* simplicity in this sense. A hypothesis that is too complicated as compared with other available hypotheses at one time can have a different status at another time if those other hypotheses have been eliminated. The first hypothesis might then be among the simplest of available hypotheses.

So, to say that the rational inquirer is concerned to find a simple hypothesis is not to say that the rational inquirer is committed to believing that “reality is simple,” whatever that might mean.

Goodman (1965) discusses a classic example. Given that all emeralds examined up until now have been found to be green, the evidence supports the hypothesis that all emeralds are green. One would not normally even consider the competing hypothesis that all emeralds are either green if first examined before A.D. 2000

or blue if not first examined before A.D. 2000. Given a suitable definition of “grue,” the latter hypothesis can be written, “All emeralds are grue.” The grue hypothesis conflicts with the first hypothesis as regards any emeralds not first examined by A.D. 2000. According to the first hypothesis those emeralds are green; according to the second they are blue.

Goodman points out the hypotheses like the second, grue hypothesis are not taken seriously. His “new riddle of induction” asks what the difference is between hypotheses like the first hypothesis that are taken seriously and hypotheses like the grue hypothesis that are not taken seriously. Clearly, there is a sense in which the answer has something to do with simplicity. The first hypothesis is much simpler than the second. But what sort of simplicity is in question and why should it be relevant?

In thinking about this it is very important to see that using simplicity to rule hypotheses out of consideration is to be distinguished from using simplicity as an explicit consideration in theory choice. Sometimes a scientist will say that a particular theory is better than another because the first theory assumes the existence of fewer objects, fewer basic principles, or whatever. When a scientist argues in some such way, he or she is arguing in favor of one rather than another hypothesis that is already being taken seriously. As Sober (1988) has observed, such appeals to simplicity are often quite controversial. That is, it is controversial whether simplicity in one or another respect is a relevant consideration in choosing among hypotheses.

But even where there are deep controversies in a subject, reasonable disputants will still take seriously only a very few of the infinitely many possible hypotheses. We are concerned with whatever it is that leads reasonable people to disregard most of the hypotheses as too “silly” to be considered (remembering that silliness is a relative matter: we can imagine circumstances in which previously ignored hypotheses become discussable).

Let us call the sort of simplicity we are concerned with “basic simplicity.” since crazy or silly hypotheses are ruled out in all domains, let us assume that there is a single domain independent notion of basic simplicity.

The basic simplicity of a hypothesis seems to have something to do with the simplicity of its representation. But it is always possible to represent any hypothesis simply, so the matter is a bit more complex. Using the new term “grue” Goodman’s second hypothesis can be represented as simply as the original hypothesis that all emeralds are green. In fact, any hypothesis can be abbreviated by a single symbol, so simplicity of surface representation cannot be taken at face value.

But suppose a hypothesis like, “All emeralds are green,” is used to explain the data. Then it has to be expanded to its more complex form, “All emeralds are either green if first examined before A.D. 2000 or blue if not first examined before A.D. 2000.” This expansion is needed on the assumption that we are more interested in accounting for the colors of objects, such as whether they are blue or green, as opposed to their “cholers,” such as whether they are grue or bleen. If instead we were more interested in explaining why emeralds were grue, we could use the hypothesis, “All emeralds are grue,” without having to expand it, and the hypothesis that “All emeralds are green,” would require elaboration in terms of grue and bleen in order to provide the desired explanation.

So, perhaps the thing to look at is not so much the mere statement of the hypothesis but also how complicated it is to use the hypothesis to explain the data and predict new observations of a sort in which we are interested. In considering possible explanations of given data, it appears to be rational and reasonable to ignore hypotheses that are much harder to use in explanation and prediction than other available hypotheses that in other respects account equally well for the data. This represents a further respect in which theoretical rationality depends on practical concerns.

6. Rationality in Action

We have seen that practical considerations play a number of roles in theoretical rationality. Practical considerations can help determine what questions it is rational to try to answer. Furthermore, practical considerations may play a role in determining what hypotheses theoretical rationality takes seriously and may also be relevant to the conservatism of theoretical reasoning. But I have not said much about practical rationality except to observe that it is often rational to make an arbitrary choice about what to do in a way that it is not rational to make an arbitrary choice of what to believe. I also noted that desires are relevant to what to decide to do in a way in which they are not relevant to what to decide to believe. If one wants something to be true, that can be a reason to try to make it true, but not by itself a reason to believe that it is true apart from what one does.

More needs to be said about how goals are relevant to practical reasoning. For example, one significant issue is whether there is a single category of goal, or a single measure of “utility,” as opposed to a variety of functionally different things: desires, values, goals, intentions, commitments, principles, rules, etc. A related issue is whether we need to allow for a structure within goals in which some goals depend on others.

Here it is not enough simply to appeal to mathematical decision theory. In the simplest decision to which the theory applies (e.g. von Neumann & Morgenstern, 1944), one is faced with a decision between two or more exclusive acts, each of which has various possible outcomes to which one assigns certain values or “utilities.” So $u(A)$ represents the utility of act A . One also assigns conditional probabilities, $prob(O,A)$, to each possible outcome O in relation to a given act A . Then the “expected gain” of a given outcome O of an act A is $(prob(O,A))u(A)$.

The “expected utility” of each act A is the sum of the expected gains of each possible consequence of that act. Then the theory holds that rationality requires

doing the act with the highest expected utility or, if there is a tie for highest, one of the acts with highest expected utility.

The principles of mathematical decision theory are like principles of logic in being principles of consistency or coherence. So it is as much of a mistake to identify mathematical decision theory with the theory of practical rationality as it is to identify the theory of theoretical rationality with logic.

Some decision theorists argue that it is useful for individuals face with hard practical problems to think of them in decision theoretic terms. Such individuals are advised to consider carefully what their possible acts are, what possible consequences each act might have, what utility they assign to each possible consequence, and how likely they think a given act would be to have a given consequence. They should then calculate expected utilities and choose the act with the highest calculated expected utility.

Is that good advice? The question can only be answered empirically. Do people do better using such a method or not? The suggested method is not obviously good advice. Given a poor enough assignment of utilities and probabilities, one can be led very wrong by one's calculation.

1.6 Derivative Goals

Some goals are derivative from others in a way that is important for practical rationality. One wants A. B is a means to A. So one wants B. That is, one wants B as a means to A. If one gets A in some other way, one no longer has the same reason to want B. Or, if one discovers that B is not going to lead to A, one no longer has the same reason to want B. It is irrational to continue to pursue an instrumental goal after the reason for wanting it has lapsed.

Also, consider the problem of deciding what to do when one has several goals. If one does A, one will satisfy goals G1, G2, and G3. If one does B, one will satisfy goals G4, G5, and G6. It is not easy to say how a rational person reaches an overall evaluation of acts A and B by combining his or her evaluation of the outcomes of each act. One idea (going back at least to Franklin, 1817) is to try

to reduce the lists by matching outcomes of A with equivalent outcomes of B, cancelling the equivalent goals out, then considering only the remaining advantages of each course of action. That can still leave difficult choices.

But one thing can be said: one should not count the satisfaction of two goals as distinct advantages of an act if one's only reason for the first is that it will enable one to attain the second! That would be to count the same consideration twice.

Another point is that one cares about some things that are neither ultimate ends nor instrumental toward getting other things one wants. One can want good news, because that is evidence for something else one wants. But the desire for good news is not a reason to try to influence what the news will be without influencing the event of which the news would be news.

1.7 Intentions

A rational person does not always reason directly from current goals, figuring out the best ways to maximize their satisfaction. That would resemble Special Foundationalism with respect to theoretical reasoning. It would ignore the role of long term intentions. Such intentions record decisions already made. Such decisions are not irrevocable, but they carry considerable weight and must not be frivolously discarded. A person incapable of maintaining long term intentions would be incapable of long term planning and would have only a low level of rationality (Bratman, 1987).

Intentions are not reducible to desires and beliefs, but put constraints on current planning of a special kind. A person's actual goals, as contrasted with things merely valued or desired, might be identified with what that person intends.

Intentions are directly related to action in ways not fully understood. Some authors think there are special intentions to do something *now*, where these are acts of will or volition, serving as immediate causes of action.

7. Conclusion

At present there is no mathematically elegant account of all aspects of reasoning. We have formal theories of implication and consistency, but these are only part of the subject. Logic, conceived as the theory of implication and consistency, and probability theory are not theories of reasoning. Conservatism, simplicity, and coherence are further features, with explanation, implication, and consistency being relevant to coherence.